MIT-IBM Watson AI Lab

Upcoming Events

Delocalized Photonic Deep Learning on the Internet's Edge

Nano Explorations
March 7, 11-11:45 a.m. ET

In this virtual talk followed by a Q&A, MIT graduate student Alexander Sludds will introduce an approach to machine learning inference based on delocalized analog processing across networks. In this approach, named Netcast, cloud-based "smart transceivers" stream weight data to edge devices, enabling ultraefficient photonic inference. Register here.

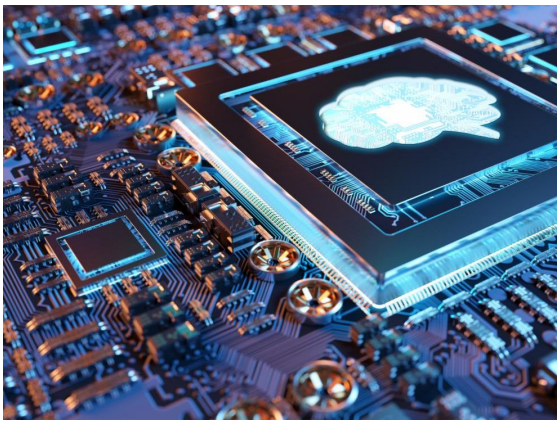

MIT AI Hardware Program Annual Symposium

MIT AI Hardware Program
April 6, 9:30 a.m.-2:30 p.m. ET

The MIT AI Hardware Program is hosting their second annual symposium. This hybrid event will include networking, poster presentations, and talks from MIT faculty, like Lab researchers co-director Aude Oliva, School of Engineering Dean Anantha Chandrakasan, Jesus del Alamo, Song Han, and Mengjia Yan, and from students. Register here.

SERC Symposium

Social and Ethical Responsibilities of Computing (SERC)
April 18, 8 a.m.-5:30 p.m. ET

Join the Social and Ethical Responsibilities of Computing for an in-person event with panel discussions including several Lab researchers on the implications of data and algorithms, beneficent and fair AI, equitable and personalized health, algorithms and humans, and ethics and computing education. Register here.

In the Lab

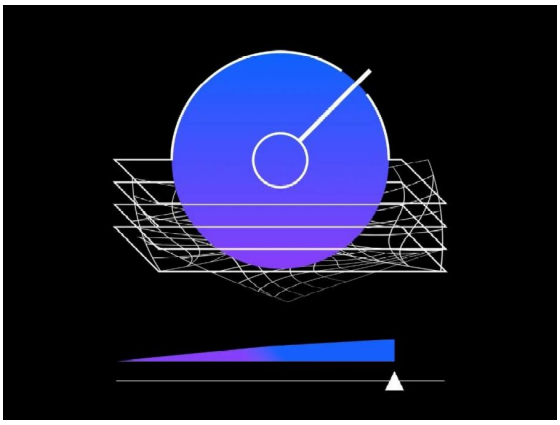

Explainer: What is prompt-tuning?

An efficient, low-cost way of adapting an AI foundation model to new downstream tasks

An emerging technique called prompt-tuning allows researcher to take off-the-shelf foundation models and customize them for new applications without retraining the model and updating its weights. Studies from the Lab explain how this method can handle multi-task learning, discover prompts for novel tasks, and be employed to help mitigate algorithmic bias and increase fairness.

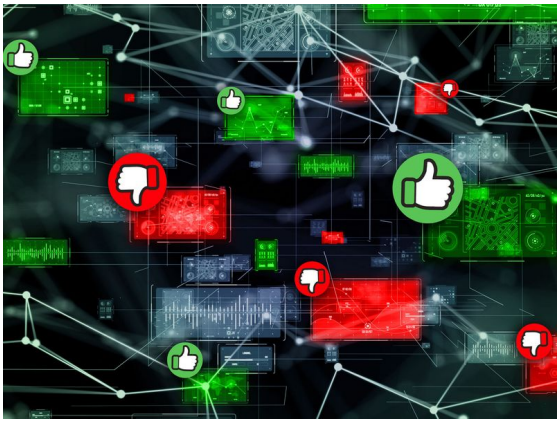

Improving machine-learning models' reliability

Efficient technique enables a model to determine its confidence in a prediction, with less compute.

A team including Lab researchers Prasanna Sattigeri, Soumya Ghosh, Subhro Das, and Gregory Wornell have developed a technique that involves creating a simpler companion model that assists the original machine-learning model in estimating uncertainty, without additional data.

Explainer: What is synthetic data?

Generated data to improve AI models, protect sensitive data, and mitigate bias

The use of generative models can provide a cost-effective way to create unlimited supply of annotated data to, for instance, fill gaps in data or supply edge cases for learning. Lab research describes ways to create tabular data, generate scenes and images that behave in a real-world manner, improve natural language processing, and test for security flaws and biases.

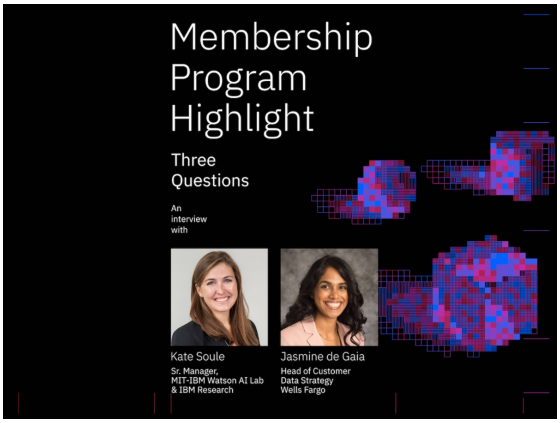

3 Questions: Innovation in financial services through synthetic data

[How generated data can preserve privacy, drive innovation and expedite technology development]

Lab's membership program lead Kate Soule sat down with Jasmine de Gaia, senior vice president and head of customer data strategy at Wells Fargo, to discuss how the Lab helped them to create synthetic test data with generative AI.

In the Media

How ChatGPT could embed a 'watermark' in the text it generates

With the rise of AI-generated text, researchers are looking at ways to identify it, including slipping in a high percentage of language that's on a "special list" flagging it to detection tools, *The New York Times* reports. But Lab co-director David Cox says new methods will likely be developed to skirt the issue. "This is always going to have an element of an arms race to it."

Foundational models and generative AI in the IBM sphere

AI has the potential and demonstrated ability to significantly impact human life from the discovery of new antibiotics to natural language processing. Here, *Forbes* shares a history of applications and IBM contributions through the work of Lab researcher Payel Das, ranging from synthetic data to the search for antimicrobials.

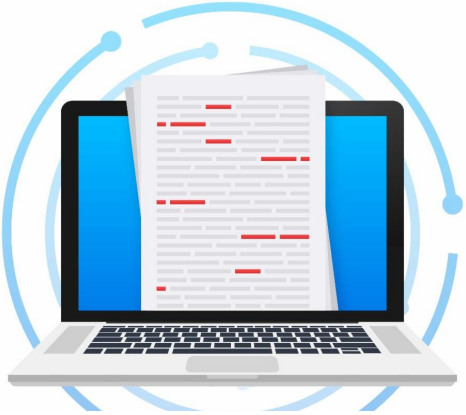

How to detect AI-generated text, according to researchers

There's more than one way to spot a fake on spotting synthetic language, *WIRED* reports. One such method from the Lab examines the likeliness of a word's usage to predict if the text was AI-generated — the higher the randomness in the text choice, the greater the chance that a human wrote it.

What business leaders need to know about quantum computing

In a piece for *CNBC*, Darío Gil, IBM Lab chair, senior vice president and director of IBM Research, lays out how quantum computing could influence business decisions and planning for the future. "I see an opportunity for businesses to step in and help higher education train tomorrow's workers, including undergrads, with the skills they'll need to help industry apply quantum computing."

AI evolving faster than even its biggest backers expected

Industry, government, and experts see the utility for AI, reports *The Sydney Morning Herald*, and so the field is evolving quickly. As such, "we'd love to harness these technologies and we're working on that for our customers, but there's an important step, which is: these things have to be trusted," says Lab co-director David Cox of the Lab's ethos.

Looking beyond microarchitectural side channels

Summing up a recent keynote talk in *ACM SIGARCH*, Lab researcher Mengjia Yan shares her perspective about the research landscape of microarchitecture security and areas worth more investigation, like the "mechanism of a state-of-the-art website-fingerprinting attack," which Lab research helped to elucidate.

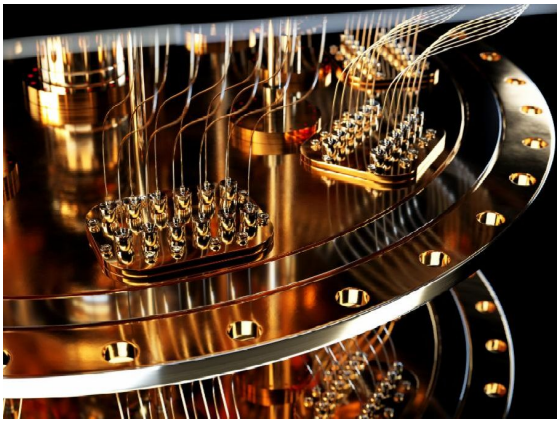

Quantum computers could solve — and create — countless problems

Quantum technology is on its way to disrupting the field of computing, which could speed up the discovery of solutions for businesses and humanitarian issues. At the same time, *Time* reports, quantum-safe protocols are increasingly needed to help protect against risks to economic and national security. "This is a new frontier of humanity," says Darío Gil, IBM Lab chair, senior vice president and director of IBM Research.

Lab Highlights

Effective July 1, Lab researchers Jacob Andreas, Kevin O'Brien, Jonathan Ragan Kelley, and Rafael Gómez-Bombarelli have been promoted to associate professor, and Elsa Olivetti will be a full professor at MIT.

Lab researchers Regina Barzilay and Markus Buehler were elected to the National Academy of

Engineering for "machine learning models that understand structures in text, molecules, and medical images" and " implementing the use of nanomechanics to model and design fracture-resistant bioinspired materials," respectively.

Lab researchers Rafael Gómez-Bombarelli, Song Han, and Jonathan Ragan-Kelley were named 2023 Sloan Research Fellows.

This summer, the Lab will host 12 IBM interns and has so far accepted 146-A undergraduate, M.Eng., and graduate students, while continuing to recruit.

Online Learning

Making AI Work: Machine Intelligence for Business and Society
A joint MIT Sloan & Schwarzman College of Computing Executive and Professional Course begins March 15.

Unsupervised Machine Learning: Unlocking the Potential of Data
A joint MIT Sloan & Schwarzman College of Computing Executive and Professional Course begins March 22.